

(10) **Patent No.:** US 9,367,321 B2  
(45) **Date of Patent:** Jun. 14, 2016

- |              |      |         |                         |         |
|--------------|------|---------|-------------------------|---------|
| 2002/0120828 | A1 * | 8/2002  | Modelski et al. ....    | 712/200 |
| 2002/0161957 | A1 * | 10/2002 | Comeau et al. ....      | 710/260 |
| 2004/0187117 | A1 * | 9/2004  | Orion et al. ....       | 718/100 |
| 2004/0194096 | A1 * | 9/2004  | Armstrong et al. ....   | 718/100 |
| 2004/0244000 | A1   | 12/2004 | Frank et al. ....       |         |
| 2005/0125795 | A1 * | 6/2005  | Kissell .....           | 718/100 |
| 2006/0059287 | A1 * | 3/2006  | Rivard et al. ....      | 710/300 |
| 2006/0095901 | A1 * | 5/2006  | Brokenshire et al. .... | 717/151 |
| 2006/0117316 | A1 * | 6/2006  | Cismas et al. ....      | 718/103 |

(Continued)

- FOREIGN PATENT DOCUMENTS

- |    |            |    |        |
|----|------------|----|--------|
| WO | 0116703    | A1 | 3/2001 |
| WO | 2006083543 | A2 | 8/2006 |

- ## OTHER PUBLICATIONS

- Andrew S. Tanenbaum, "Structured Computer Organization", 1984, Second Edition, pp. 10-12.\*


(Continued)

- Primary Examiner — Idriss N Alrobaye  
Assistant Examiner — Keith Nielsen  
(74) Attorney, Agent, or Firm — Sughrue Mion, PLLC

- (57) **ABSTRACT**

- The invention provides a processor comprising: an execution unit, and a thread scheduler configured to schedule a plurality of threads for execution by the execution unit in dependence on a respective runnable status for each thread. The execution unit is configured to execute thread scheduling instructions which manage the runnable statuses. The thread scheduling instructions including at least: one or more source event enable instructions each of which sets an event source to a mode in which it generates an event dependent on activity occurring at that source, and a wait instruction which sets one of said runnable statuses to suspended pending one of the events upon which continued execution of the respective thread depends. The continued execution comprises retrieval of a continuation point vector for the respective thread.

- 48 Claims, 9 Drawing Sheets**



(56)

**References Cited**

U.S. PATENT DOCUMENTS

2006/0136915	A1 *	6/2006	Aingaran et al. ....	718/100
2006/0179194	A1 *	8/2006	Jensen .....	710/111
2006/0195683	A1 *	8/2006	Kissell .....	712/228
2007/0124736	A1 *	5/2007	Gabor et al. ....	718/108
2007/0180438	A1 *	8/2007	Suba .....	717/151

2008/0104296	A1 *	5/2008	Blackmore et al. ....	710/267
2008/0115113	A1 *	5/2008	Codrescu et al. ....	717/127

OTHER PUBLICATIONS

International Search Report and Written Opinion, dated Jun. 24, 2008.

\* cited by examiner

Figure 1

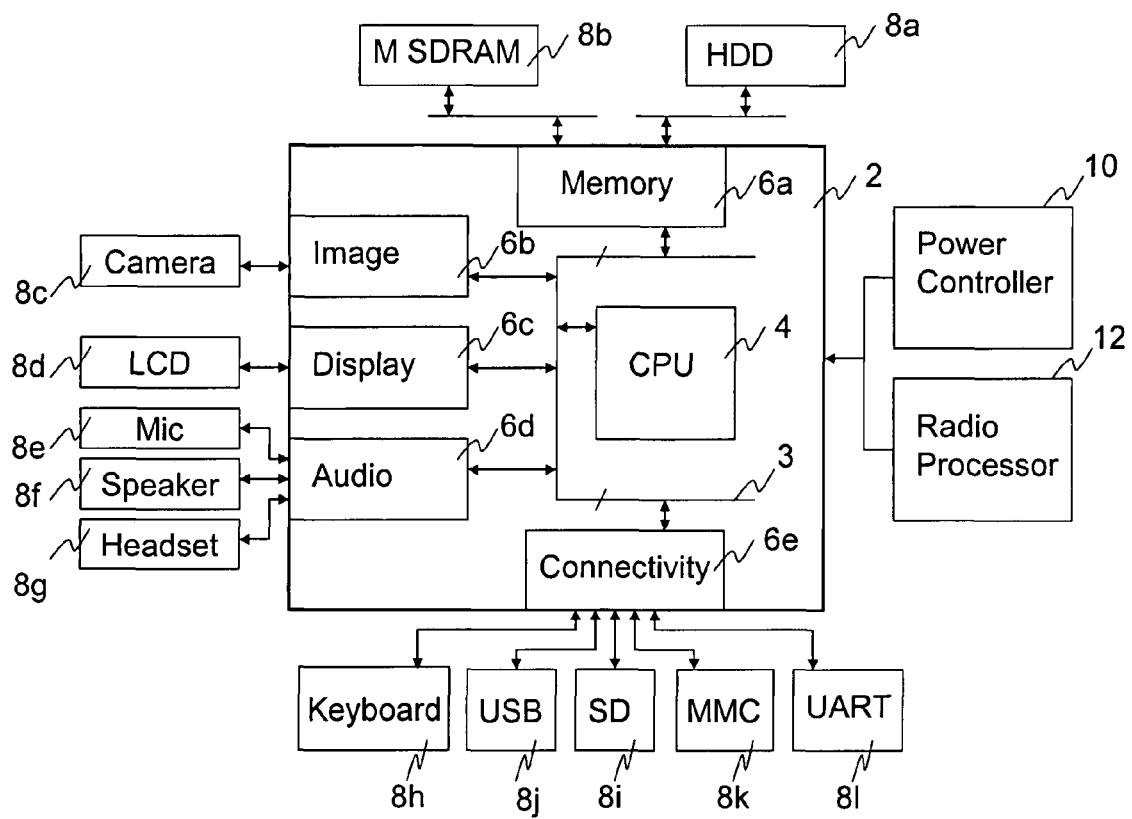


Figure 2




Figure 2A




Figure 3




Figure 4




Figure 5






Figure 6

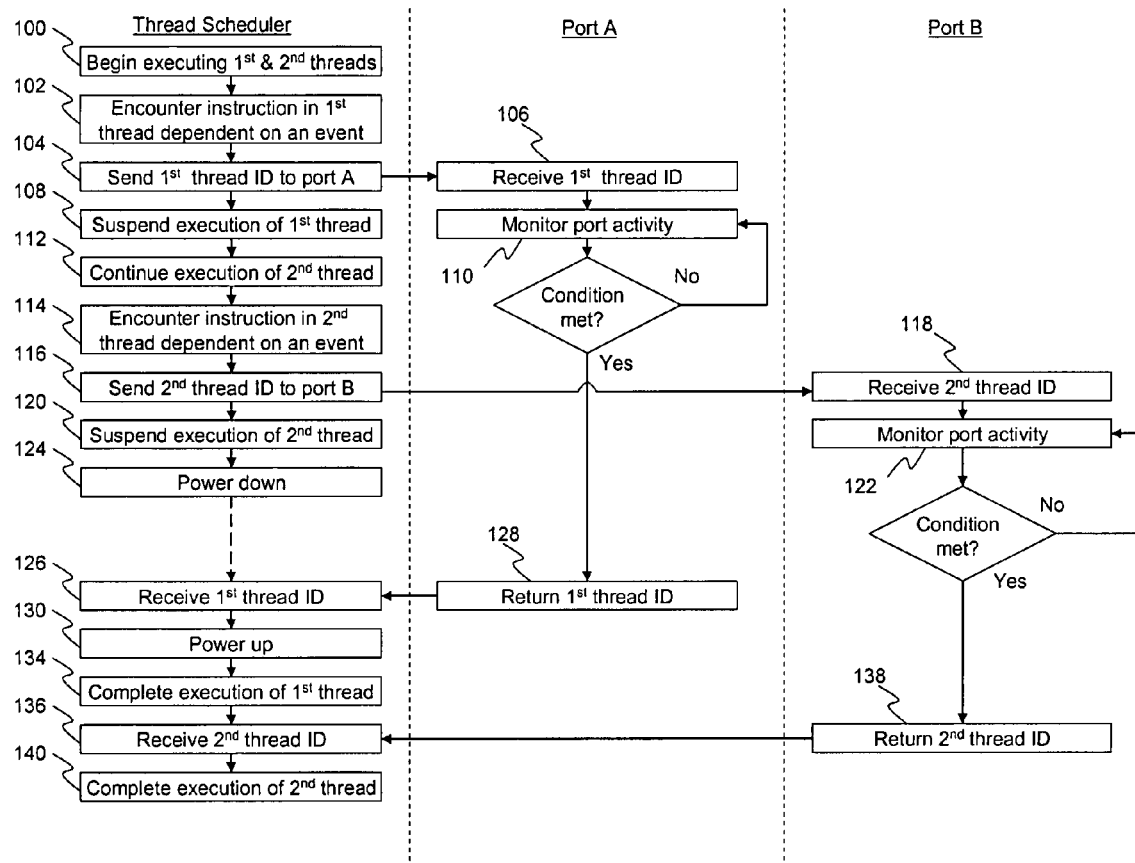


Figure 7

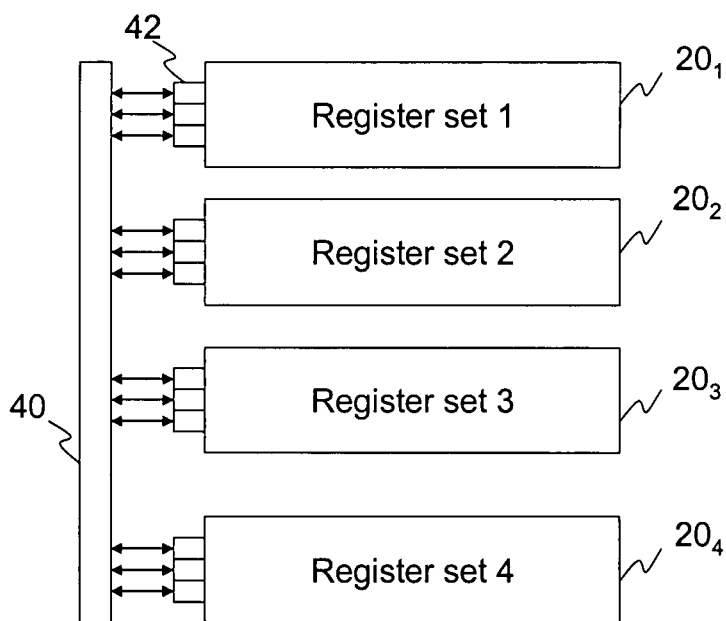



Figure 7A



1

# PROCESSOR INSTRUCTION SET FOR CONTROLLING AN EVENT SOURCE TO GENERATE EVENTS USED TO SCHEDULE THREADS

## FIELD OF THE INVENTION

The present invention relates to a processor instruction set, and particularly but not exclusively an instruction set for scheduling instructions based on certain types of activity.

## BACKGROUND OF THE INVENTION

One of the challenges facing processor designers is the handling of an ever-increasing number of communications, both internal and external to the processor. Generally this is done by providing some kind of interrupt handling capability for the processor for reacting to certain types of activity. Increasingly, more sophisticated interface logic is used to deal with, for example, multiple external devices per port.

Such capability is needed in a wide variety of different contexts. One context which is discussed herein by way of a background example is in mobile applications processing.

FIG. 1 shows an exemplary application of a mobile applications processor 2. The applications processor 2 comprises a CPU 4 and a plurality of interface controllers 6 which interface with a plurality of peripheral devices 8. The interface controllers include: a memory controller 6a for interfacing with a hard-drive (HDD) 8a and a SDRAM memory 8b; a video controller 6b for interfacing with a camera 8c; a display controller 6c for interfacing with an LCD display 8d; an audio controller 6d for interfacing with a microphone 8e, speaker 8f and headset 8g; and a connectivity controller 6e for interfacing with a keyboard 8h, a Universal Serial Bus (USB) device 8i, a Secure Digital (SD) card 8j, a Multi-Media Card (MMC) 8k, and a Universal Asynchronous Receiver/Transmitter (UART) device 8l. The interface controllers 6 are typically connected to the CPU 4 via a bus 3. The system also comprises a power controller 10 and radio processor 12.

Note that the interface controllers 6 are shown somewhat schematically, but represent generally some kind of dedicated I/O logic or specially configured ports.

Conventionally, external interfacing is achieved either using interrupts or by polling. When interrupts are used, an external peripheral device sends a signal to inform the processor either that it has data ready to input to the processor or that it requires data from the processor. However, using interrupts, the current program state must be saved before the interrupt can be acted upon. When polling is used, the processor continually checks the state of the device to determine whether or not it is ready to supply or accept data. This introduces a delayed reaction time. Polling is also slow because of the continual queries and responses.

One possibility for implementing an applications processor 2 such as that of FIG. 1 is to use an Application Specific Integrated Circuit microcontroller (ASIC). ASICs are hard-wired devices possibly including microprocessors dedicated to a particular application and optimised to suit that application. For a given function, they are generally cheaper and consume less power than other options. However, they are complex to design, must be pre-designed and cannot readily be reconfigured.

Another possibility is to use Field Programmable Gate Array (FPGA) devices. FPGAs are semiconductor devices that can be configured “in the field” after manufacture. To configure an FPGA, first a computer is used to model the desired logical functions, for example by drawing a sche-

2

matic diagram or creating a text file describing the functions. The FPGA comprises an array of look-up tables which communicate via statically configured interconnects. The computer model is compiled using software provided by the FPGA vendor, which creates a binary file that can be downloaded into the FPGA look-up tables. This allows manufacturers of equipment to tailor the FPGA to meet their own individual needs.

In this example, the interface controllers 6 are implemented as FPGAs. This has the benefit that the manufacturer of the mobile telephone can purchase generic FPGA devices 2 and then configure them on site (i.e. “in the field”) to be specific to their desired application. The disadvantage of FPGAs however is that they are more expensive, slower and consume more power than ASICs.

In alternative examples, the whole chip 2 could be implemented in FPGA, or the chip 2 could be a general purpose processor with separate FPGA chips connected between the chip 2 and the respective peripherals 8. However, these options would be even more expensive and power-consuming—prohibitively so for most mobile phones and other consumer devices.

It would be advantageous to achieve the configurability of an FPGA but with the price, speed, scope and energy consumption levels of an ASIC.

## SUMMARY OF THE INVENTION

In order to add new functionality to a processor, it must be configured to recognise new instructions, or op-codes, and to act upon those instructions in the desired manner.

Accordingly, in one aspect of the present invention there is provided a processor comprising: an execution unit; and a thread scheduler configured to schedule a plurality of threads for execution by the execution unit in dependence on a respective runnable status for each thread; wherein the execution unit is configured to execute thread scheduling instructions which manage said runnable statuses, the thread scheduling instructions including at least: one or more source event enable instructions each of which sets an event source to a mode in which it generates an event dependent on activity occurring at that source, and a wait instruction which sets one of said runnable statuses to suspend the respective thread until generation of one of said events upon which continued execution of that thread depends; wherein said continued execution comprises retrieval of a continuation point vector for the respective thread.

These new instructions and associated functionality advantageously allow the processor to be “primed” to respond quickly to events. Using suspended threads, the thread scheduler can prepare to execute a thread in expectance of an event, thus enabling this fast response time. In contrast, using conventional interrupts, the execution unit is interrupted by a signal whilst executing some potentially unrelated code, and without having made any preparations for acting upon the interrupt.

Also, using a wait instruction and separate event enable instruction decouples the operation of setting up the event from the operation actually suspending the thread. For example, in embodiments, this advantageously allows multiple events to be set up for a given thread by enabling or disabling different event sources using multiple event source enable instructions, and the thread can then be suspended pending activity from any one of those enabled sources using a single wait instruction. That is, the execution unit may be adapted to execute a plurality of said event source enable instructions each of which sets a respective event source to a

mode in which it generates an event dependent on activity occurring at that source, and said wait instruction may set said one of the runnable statuses to suspend the respective thread until generation of an event from any of said event-enabled sources.

The thread scheduling instructions may include a source event disable instruction which sets an event source to a mode in which it does not generate the event.

The thread scheduling instructions may include: a thread event enable instruction which sets a control status of a thread to event-enabled to allow the enabled thread to accept events, and a thread event disable instruction which sets a control status of a thread to event-disabled to stop the disabled thread from accepting events.

The execution unit may be configured to execute one or more additional instructions each of which configures a parameter associated with an event source.

The additional instructions may include a set condition instruction which sets a condition upon which one of said event sources generates the respective event. The additional instructions may include a set data instruction which provides data associated with the condition.

The continuation point vector may be one of said parameters associated with one of said event sources, and the additional instructions may include at least a set vector instruction which sets said continuation point vector.

The decoupling of the event set-up from the wait instruction may also advantageously allow events to be left configured at sources event when the event is disabled. That is, the processor may be arranged to execute a first one of said event source enable instructions to enable a first event source to generate a first event dependent on a first activity occurring at the first source, to execute at least one of said additional instructions to configure one or more of said parameters for the first event source, to execute said wait instruction to suspend the respective thread pending the first activity, to disable the first source from generating the first event after occurrence of the first activity, to leave said parameters configured for the first event source after said disabling of the first source, and to re-use said parameters by executing a second of said event source enable instructions to enable the first source to generate a further instance of the first event dependent on a further occurrence of the first activity at the first source.

The thread scheduling instructions may further include an interrupt enable instruction which sets a control status of a thread to interrupt-enabled to allow the enabled thread to accept interrupts, and an interrupt disable instruction which sets a control status of a thread to interrupt-disabled to stop the disabled thread from accepting interrupts.

The thread scheduling instructions may further include: an input instruction which pauses a thread pending input of data from an event source, and an output instruction which pauses a thread pending the availability of an event source for outputting data; wherein continued execution of a thread paused by said input and output instruction does not involve retrieval of a continuation point vector for that thread.

The processor may be arranged to execute at least one of said input and output instructions between said execution of said wait instruction and said execution of said second event source enable instruction.

The thread scheduling instructions may include a clear instruction to disable all events for a thread.

The processor may comprise at least one first register for storing said control statuses. The processor may comprise at least one second register for storing said parameters.

The processor may comprise an interconnect system for establishing at least one channel between at least two third

registers each arranged to store information relating to respective thread, and at least one channel may be an event source.

The processor may comprise at least one port being an event source.

Said wait instruction may be able to be either a wait enable true instruction which waits only if its condition operand is true, or a wait enable false instruction which waits only if its condition operand is false. Said event source enable instruction may be able to be either a source event enable true instruction which enables the source to generate events if an operand is true and disables it otherwise, or an source event enable false instruction which enables a source to generate events if an operand is false and disables it otherwise.

The execution unit may be arranged to execute one of said thread event enable instructions and subsequently execute a plurality of said source event enable instructions each for a respective source.

The processor may be adapted to automatically disable a thread from accepting events upon occurrence of the respective event without executing a thread event disable instruction. The processor may be adapted to automatically disable the event source from generating events upon occurrence of a respective event without executing a source event disable instruction.

The processor may be adapted to complete at least one wait, input or output instruction immediately if the respective event, data or availability is ready on or before execution of the respective instruction.

According to another aspect of the invention, there is provided a method of controlling a thread scheduler to schedule threads for execution by an execution unit within a processor, the method comprising: scheduling a plurality of threads in dependence on a respective status for each thread; and operating the execution unit to execute thread scheduling instructions for managing statuses of threads, said thread scheduling instructions including at least: one or more source event enable instructions each of which sets an event source to a mode in which it generates an event dependent on activity occurring at that source, and a wait instruction which sets one of said runnable statuses to suspend the respective thread until generation of one of said events upon which continued execution of that thread depends; wherein said continued execution comprises retrieval of a continuation point vector for the respective thread.

According to another aspect of the invention, there is provided an execution unit configured to execute thread scheduling instructions which manage statuses of threads, the thread scheduling instructions including at least: one or more source event enable instructions each of which sets an event source to a mode in which it generates an event dependent on activity occurring at that source, and a wait instruction which sets one of said runnable statuses to suspend the respective thread until generation of one of said events upon which continued execution of that thread depends; wherein said continued execution comprises retrieval of a continuation point vector for the respective thread.

According to another aspect of the invention, there is provided a method of scheduling a plurality of threads for execution by an execution unit, the method comprising executing thread scheduling instructions for managing statuses of threads, said thread scheduling instructions including at least: one or more source event enable instructions each of which sets an event source to a mode in which it generates an event dependent on activity occurring at that source, and a wait instruction which sets one of said runnable statuses to suspend the respective thread until generation of one of said

events upon which continued execution of that thread depends; wherein said continued execution comprises retrieval of a continuation point vector for the respective thread.

For a better understanding of the present invention and to show how the same may be carried into effect, reference will now be made, by way of example, to the corresponding drawings.

#### DETAILED DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an example application of an FPGA device as is known in the art;

FIG. 2 illustrates an example application of an interface processor;

FIG. 2a illustrates another example application of an interface processor;

FIG. 3 is a schematic representation of the architecture of an interface processor;

FIG. 4 is a schematic representation of a port;

FIG. 5 is a schematic representation of thread register sets;

FIG. 6 is a flow diagram illustrating the operation of the thread scheduler;

FIG. 7 is a schematic representation of channels between thread register sets; and

FIG. 7A is a schematic representation of a channel end.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

FIG. 2 shows an exemplary application of interface processors in a mobile telephone. The mobile applications processor 2 needs to communicate with the plurality of peripheral devices 8. As in FIG. 1, the applications processor 2 comprises a bus 3, a CPU 4, and a memory controller 6a for interfacing with a hard-drive (HDD) 8a and a SDRAM memory 8b, as well as a power controller 10 and radio processor 12.

However, in place of dedicated controllers 6, the arrangement of FIG. 2 allows the CPU 4 to communicate externally via generic ports 7. In this example, generic ports 7a and 7b are provided for interfacing with cameras 8c and LCD displays 8d; a generic port 7c is provided for interfacing with a microphone 8e, speaker 8f and headset 8g; and a generic port 7d is provided for interfacing with a keyboard 8h, a Universal Serial Bus (USB) device 8i, a Secure Digital (SD) card 8j, an Multi-Media Card (MMC) 8k, and a Universal Asynchronous Receiver/Transmitter (UART) device 8l.

In FIG. 2, interface processors 14a, 14b and 14c are placed at the outputs of the relevant ports 7, with a first interface processor 14a being connected between the image devices 8c-8d and the generic ports 7a-7b, a second interface processor 14b being connected between the audio devices 8e-8g, and a third interface processor 14b being connected between the generic port 7d and the various connectivity devices 8h-8m. The ports 7 need only be general purpose ports because the application-specific display, audio and connectivity functions are implemented by the interface processors 14a-14c in a manner to be described later. The ports 7 need not use FPGA logic, because the interface processors 14 provide the flexibility and configurability that would otherwise be provided by FPGAs. The interface processor 14a has ports 22a and 22b connected to the ports 7a and 7b, and ports 22c, 22d, 22e and 22f connected to the external devices 8c and 8g. The interface processors 14b and 14c have similar ports, not shown in FIG. 2.

The interface processors are typically involved in implementing the specific protocols used to transfer data via the interfaces, re-formatting data including converting it between parallel and serial formats, and possibly higher level functions such as encoding it, compressing it or encrypting it.

Another application of an interface processor is as a tile in a multiprocessor chip 202 illustrated in FIG. 2a. Such a chip 202 uses a high performance interconnect 204 which supports communication between the processors 14 on the chip 202, and inter-chip links 206 so that systems can easily be constructed from multiple chips. Each processor 14 is connected to its inter-chip link 206 and high performance interconnects 204 via ports 22.

An important feature of the interface processor which is discussed more fully in the following is its ability to manage activity at the ports 22. Each interface processor comprises a CPU, memory and communications. To allow the direct and responsive connectivity between the CPU and the ports, each processor has hardware support for executing a number of concurrent program threads, each comprising a sequence of instructions, and at least some of which are specifically responsible for handling activity at the ports. As will be discussed more fully in the following, the hardware support includes:

- a set of registers for each thread,
- a thread scheduler which dynamically selects which thread to execute,
- a set of ports used for input and output (ports 22),

The use of a small set of threads on each processor can be used to allow communications or input/output to progress together with other pending tasks handled by the processor, and to allow latency hiding in the interconnect by allowing some threads to continue whilst others are suspended pending communication to or from remote interface processors.

FIG. 3 shows schematically an exemplary architecture of an interface processor 14 according to one embodiment of the present invention. The processor 14 comprises an execution unit 16 for executing threads of instructions under the control of a thread scheduler 18. The processor 14 further comprises a Random Access Memory (RAM) 24 for holding program code and other data, and a Read Only Memory (ROM) (not shown) for storing permanent information such as boot code.

The thread scheduler 18 dynamically selects which thread the execution unit 16 should execute. Conventionally, the function of a thread scheduler would simply be to schedule threads from the program memory in order to keep the processor fully occupied. However, according to the present invention, the scheduling by the thread scheduler 18 is also related to activity at the ports 22. It is noted in this respect that the thread scheduler may be directly coupled to the ports 22 so as to minimise the delay when a thread becomes runnable as a result of an input or output activity at the port.

Each of the m threads under consideration by the thread scheduler 18 is represented by a respective set of thread registers 20<sub>1</sub> . . . 20<sub>m</sub> in a bank of registers 20, to which the thread scheduler 18 has access. Instruction buffers (INSTR) 19 are also provided for temporarily holding instructions fetched from memory 24 before being subsequently issued into the execution unit 16. The details of these registers and buffers are discussed later.

Of the m threads, the thread scheduler 18 maintains a set of n runnable threads, the set being termed "run", from which it takes instructions in turn, preferably in a round-robin manner. When a thread is unable to continue it is suspended by removing it from the run set. The reason for this may be, for example, because the thread is awaiting one or more of the following types of activity:

7

its registers are being initialised prior to it being able to run, it has attempted an input from a port which is not ready or has no data available,  
it has attempted an output to port which is not ready or has no room for the data,  
it has executed an instruction causing it to wait for one or more events which may be generated when ports become ready for input.

Note that the term “event” as used herein refers to a particular type of operation, which is slightly different from basic input-output operation. The distinction is discussed below in relation to FIGS. 4 and 5.

Advantageously, in order to facilitate rapid reaction time, a direct hardwired connection 28 is provided between the thread scheduler 18 and the execution unit 16 to allow the thread scheduler 18 to control which thread or threads the execution unit 16 should fetch and execute. Direct hardwired paths 30a, 30b, 30c are also provided between the thread scheduler 18 and each of the ports 22; and direct hardwired paths 29<sub>1</sub> . . . 29<sub>m</sub> are provided between the thread scheduler 18 and each of the registers 20. These direct paths preferably provide control paths which allow the thread scheduler to associate a respective thread with one or more of the ports 22, and particularly to return ready indications from the ports when certain activity occurs, allowing the processor to respond quickly to activity or stimuli occurring at the ports 22. The operation of the thread scheduler in relation to the ports is discussed below with regard to FIGS. 4 and 6.

The execution unit 16 also has access to each of the ports 22a-22c and each of the registers 20<sub>1</sub>-20<sub>m</sub> via direct connections 27 and 31, thus providing a direct link between the core processor, registers, and the external environment. Preferably, these direct paths provide further control paths allowing the execution unit to pass conditions to the ports. This is discussed in further detail below with regard to FIG. 4. The direct paths 27 and 31 may also allow data to be input and output directly between the thread registers 20 and the ports 22, thus allowing threads to communicate directly with the external environment. For example, data may be written directly from an external device to one of a thread's operand registers, rather than being written to memory 24 and then subsequently fetched. Conversely, following an operation, data from an operand register may be picked up by the execution unit 16 and sent directly out of a port 22. This improves reaction time significantly.

Note that by “direct connection” or “direct path” it is meant a connection separate from the connection between the execution unit and the program memory 24. Thus, for example, the thread scheduler 18 and execution unit 16 have access to data input from ports 22 without that data being stored and then subsequently fetched from memory 24. Particularly, if the connection between the execution unit 16 and memory 24 is via a bus 3, then a “direct” connection or path means one which is separate from the bus. Thus the various communications between ports 22, registers 20, thread scheduler 18 and execution unit 16 can all occur without the need for bus arbitration, improving reaction time. The ports 22 may also be provided with an additional connection (not shown) with the bus 13.

The term “port” as used in this application can refer to either a “pin port” or a “data port”. A pin port is responsible for detecting individual logical transitions, i.e. rising and falling edges, of a signal occurring at a pin at the processor chip's physical boundary. Data ports are “higher level” in that they can handle one or more bits, typically accumulated in an I/O buffer, and typically making up a portion of data such as a word. Instead of detecting rising and falling edges, a data

8

port handles the state or logic level of a bit or bits at a particular instant. A data port may be on/off chip, or it may be a port to another processor embedded on the same chip. Note that “pin port” and “data port” may in fact refer to different modes of the same actual port.

FIG. 4 shows schematically a port 22 according to a preferred embodiment of the invention. The port 22 comprises an I/O buffer 32 for passing input and output data to and from the processor 14. In addition, each port 22 comprises activity handling logic 36 for monitoring activity occurring at the port and signalling the occurrence of certain activity by means of at least one ready bit or flag 37. The ready flag 37 is preferably signalled to the thread scheduler via direct path 30. Potential activity which the port may detect includes:

- data has been input to the port,
- some specific data has been input to the port, and/or
- the port has become available for output.

To facilitate the detection of such activity, the port 22 is provided with a set of registers 38. These comprises a thread identifier (TID) register for storing an identification of the relevant thread, a control (CTRL) register for storing one or more conditions, a continuation point vector (VECTOR) register for storing the position in the program where execution was suspended, and a data (DATA) register for storing any data associated with a condition. The values TID is written to the registers 38 by the thread scheduler 18 via the direct path 30 (which would be 30a, 30b, 30c in FIG. 3), and the values VECTOR, CTRL and DATA are written by the execution unit 16 via the direct path 31. The TID is returned to the thread scheduler 18 upon detection of the desired activity in order to identify the associated thread. The activity logic also comprises an enable flag 39, which is discussed in further detail below.

Note that although the registers 38 are shown in FIG. 4 as being contained within the port 22, they may in fact be situated elsewhere within the processor 14 and simply associated with the port 22.

FIG. 5 shows an exemplary bank of thread registers 20 used to represent the threads. The bank 20 comprises a plurality of sets of registers corresponding to respective threads T<sub>1</sub> to T<sub>m</sub> which are currently under consideration by the thread scheduler 18. In this preferred example, the state of each thread is represented by eighteen registers: two control registers, four access and twelve operand registers. These are as follows.

Control Registers:

PC is the program counter

SR is the status register

Access Registers:

GP is the global pool pointer

DP is the data pointer

SP is the stack pointer

LR is the link register

Operand Registers: OP1 . . . OP12

The control registers store information on the status of the thread and for use in controlling execution of the thread. Particularly, the ability of a thread to accept events or interrupts is controlled by information held in the thread status register SR. The access registers include a stack pointer used for local variables of procedures, a data pointer normally used for data shared between procedures and a constant pool pointer used to access large constants and procedure entry points. The operand registers OP1 . . . OP12 are used by instructions which perform arithmetic and logical operations, access data structures, and call subroutines.

A number of instruction buffers (INSTR) 19 are also provided for temporarily storing the actual instructions of the thread. Each instruction buffer is preferably sixty-four bits

long, with each instruction preferably being sixteen bits long, allowing for four instructions per buffer. Instructions are fetched from program memory **24** under control of the thread scheduler **18** and placed temporarily in the instruction buffers **19**.

The execution unit has access to each of the registers **20** and buffers **19**. Further, the thread scheduler **18** has access to at least the status register SR for each thread.

As mentioned above, the term “event” as used herein refers to a particular type of operation, or to the activity corresponding to that particular type of operation. Event based operations are slightly different from basic input-output operations, and work as follows. An event is first set for a thread by transferring a continuation point vector from the execution unit **16** and a thread identifier from the thread scheduler **18** to the VECTOR and TID registers **38** associated with a port **22**, preferably via direct paths **31** and **30**. An associated condition and condition data may also be written to the CTRL and DATA registers **38** of the port **22**. The event is thus set at the port, but not necessarily enabled. To enable the port to generate an indication of an event, the port’s enable flag **39** must also be asserted, preferably by the thread scheduler **18** via direct path **30**. Further, to enable the thread itself to accept events, the thread’s event enable (EE) flag in the respective status register SR for the thread must be set to event-enabled. Once the event is thus set and enabled, the thread can be suspended awaiting the event using an event-based wait instruction which acts on the thread scheduler **18**. At this point, the current pending instruction may be discarded from the relevant instruction buffer **19**. When the event occurs, e.g. some data is input to the port, the occurrence is signaled by the return of the thread identifier and continuation point vector from the port **22** to the thread scheduler **18** and execution unit **16**, allowing the instruction at the continuation point vector to be fetched from program memory **24** into an instruction buffer **19** and execution resumed at the appropriate point in the code.

When the event occurs, the thread’s EE flag in the respective status register SR may be set to event-disabled to prevent the thread from reacting to events immediately after the event occurs. The enable flag **39** may be de-asserted as a result of the thread executing instructions when the event occurs.

The enable flag **39** can be asserted whilst setting up a number of ports in preparation for waiting for an event from one or more of the ports. The thread’s EE flag may also be set to event-enabled prior to enabling a set of port enable flags and in this case the first port to be enabled which is ready will generate an event causing the current instruction to be discarded and execution to proceed by immediately fetching and executing the instruction at the continuation point vector.

The advantage of the port’s enabling flag **39** and status register EE flag is that the enabling and disabling of events is separated from both the setting up of the events and the suspension of a thread by a wait instruction, allowing different input and output conditions to be readily toggled on and off for a particular thread and/or for various different threads. For example, an event may be left set up at a port **22** even though the event is disabled. Thus events may be re-used by a thread because, although the event has already occurred once, the thread identifier, continuation point vector and condition are still stored in the TID, VECTOR, CTRL and DATA registers **38** of the port **22**. So if the thread needs to re-use the event, the port’s registers **38** do not need to be re-written, but instead the port’s enable flag **39** can simply be re-asserted and/or the EE flag in the status register SR for a thread can be re-set to event-enabled. A further wait instruction will then suspend the thread pending a re-occurrence of the same event.

Furthermore, the use of continuation point vectors allows multiple events to be enabled per thread. That is, a given thread can set up one event at one port **22a** by transferring a continuation point vector to that port, set up another event at another port **22b** by transferring a different continuation point vector to that other port, and so forth. The thread can also enable and disable the various events individually by separately asserting or de-asserting the different enable flags **39** for each respective port. A wait instruction will then cause the thread to be suspended awaiting any enabled event.

In contrast with events, using basic I/O operations the thread scheduler **18** does not transmit a continuation point vector to the VECTOR register, and does not use the port’s enable flag **39** or the EE flag in the status register SR. Instead, the pending instruction is simply left in an instruction buffer **19**, and if necessary execution is simply paused pending either an input or the availability of the port for output, as indicated by the ready flag **37**. In embodiments, only the TID register may be required for scheduling according to a basic I/O. A basic I/O may or may not use a condition in the CTRL and DATA registers. If such a condition is not used, the I/O will simply be completed as soon as the port is ready.

Note also that once execution of a thread is resumed following an event, it may of course subsequently perform a basic I/O operation. Conversely, once a thread is resumed following a basic I/O, it may subsequently include an event operation. Any such chain of events and I/Os may be included in a thread. For example, a basic I/O operation may be interleaved between two event-based wait operations while the event is disabled (i.e. while the port’s enable flag **39** and/or the status register’s EE flag is de-asserted) but while the event vector and condition are still left set in the registers **38**. That is, the event may be disabled following completion of a first event-based wait operation, a basic I/O subsequently performed using the same port, and then the same event re-enabled for use in a second event-based wait operation. As discussed above, the basic I/O operation pauses and un-pauses the thread but does not effect the port’s enable flag **39** or the EE flag in the status register, nor transfer control to the event vector.

The operation of the thread scheduler and two exemplary ports is now described with reference to the flow diagram of FIG. 6. In step **100** the execution unit **16** begins execution of a first and a second thread under the direction of the thread scheduler **18**. At step **102** the thread scheduler encounters a portion of code within the first thread which is conditional on an event, for example the receipt of some particular signal at the port **22a**. At step **104**, the thread scheduler sends the thread identifier (TID) of the first thread to the Port **22a**, along with a continuation point vector (VECTOR) specifying at what point in the program thread execution should be continued once the event is detected, and also any required condition control information (CTRL) and associated data (DATA). For example, the data may be the value of the signal which the instruction is expecting to receive at the port. At step **104**, the thread scheduler may also set the enable flag **39** of the first port and set the status register SR of the first thread to event-enabled.

At step **106** the port **22a** receives this information from the thread scheduler **18**. At step **108** the thread scheduler **18** suspends execution of the first thread. At step **110** the port **22a** begins to monitor the activity at that port.

At step **112** the thread scheduler **18** determines that the second thread is still outstanding and the execution unit **16** continues execution of the second thread under the direction of the thread scheduler **18**. In step **114** the thread scheduler **18** encounters a portion of code which is conditional on an event.



## 11

At step 116 the thread scheduler 18 sends the thread identifier, along with the continuation point vector and any other required condition information, to the port 22b. At step 116, the thread scheduler may also set the enable flag 39 of the second port and set the second status register for the second thread to event-enabled. At step 118 the port 22b receives this information. At step 120 the thread scheduler suspends execution of the second thread. At step 122 the port 22b begins to monitor the activity occurring at that port.

At step 124 the thread scheduler determines that there are currently no more outstanding threads to be scheduled and the system powers down all components except for the ports 22a and 22b. At step 128 the port 22a detects the relevant event, for example the receipt of the signal stored in the DATA register, and consequently returns the thread identifier (TID) and continuation point vector (VECTOR) (as well as setting the status register of the first thread to event-disabled). At step 126 the thread scheduler 18 receives the returned identifier. Now that execution can continue, at step 130 the system powers up again. At step 134 the execution unit 16 completes the execution of the first thread under the direction of the thread scheduler 18. At step 138 the port 22b detects the relevant event for the second thread and returns its thread identifier and continuation point vector (as well as setting the status register of the second thread to event-disabled). At step 136 the thread scheduler 18 receives the returned information, and at step 138 the execution unit 16 completes the execution of the second thread under the control of the thread scheduler 18. Note that there could be an additional powering down step between steps 134 and 136.

As illustrated in FIG. 7, the principle of the invention can also be extended to communication between threads, or more accurately between the thread register sets 20 which store information representing the threads. For illustrative purposes, only four thread register sets 20<sub>1</sub> to 20<sub>4</sub> are shown in FIG. 7, each storing information for a respective thread T<sub>1</sub> to T<sub>4</sub>. Each of the thread register sets is connected to each of the other sets by an interconnect system 40, which is operable to establish at least one channel for transferring data directly between at least two of the thread register sets 20. Channels are preferably used to transfer data to and from the operand registers OP, but could in principle be used to transfer information to or from other types of register such as a status register SR. The thread scheduler 18 can schedule threads based on activity occurring over channels in a similar manner as discussed in relation to ports above.

As shown in FIG. 7A, each of the channel ends 42 resembles a pair of ports, with an input buffer 44 and an output buffer 46. Similarly to the ports 22, each channel input buffer 44 and output buffer 46 may have activity handling logic 36' for monitoring activity occurring over a channel and signalling the occurrence of certain activity by means of at least one ready flag 37'. Potential activity which the channel end may detect includes: that data has been input to the channel, or that the channel has become available for output. If an output instruction is executed when the channel is too full to take the data then the thread which executed the instruction is paused, and is restarted when there is enough room in the channel for the instruction to successfully complete. Likewise, when an instruction is executed and there is not enough data available then the thread is paused, and will be restarted when enough data becomes available.

Again as with the ports 22, in order to facilitate the detection of such activity, each channel end is associated with registers 38'. These comprise a thread identifier (TID) register for storing an identification of the relevant thread, and a continuation point vector (VECTOR) register for storing the

## 12

position in the program where execution should resume upon occurrence of an event. These TID and VECTOR registers can then be used by the thread scheduler 18 and execution unit 16 to schedule threads in the same manner as with the ports 22. The VECTOR register allows the channel to generate events and interrupts. The channel end also has an enable flag 39' to enable the channel to generate events. In embodiments, the channel ends 42 may not be provided with CTRL and DATA registers.

The same channel ends 42 may also be used to communicate data from the thread registers to the external environment via the ports 22. That is, the execution unit 16 may pick up the contents of a register 20 via a channel end 42 and pass it directly out via a port 22; and conversely, the execution unit 16 may also receive input from a port 22 and transfer it directly to a register 20 via a channel end 42. Thus if two or more interface processors according to the present invention are connected together, as shown for example in FIG. 2A, a channel can also be established between those processors. Within an inter-processor system, each processor may interface with a system interconnect via a number of bidirectional physical links 204 which allow simultaneous connection with other processors.

The general term used herein to cover ports, channels, and other sources of activity is "resource".

The interface processor can support several programming approaches due to its thread-based structure. It can be treated as a single conventional processor performing standard input and output, or it can be programmed as part of a parallel array of hundreds of communicating components. An instruction set is provided which supports these options. The instruction set includes special instructions which support initialisation, termination, starting and stopping threads and provide input/output communication. The input and output instructions allow very fast communications with external devices. They support high-speed, low-latency input and output and high-level concurrent programming techniques. Their application therein to handling port activity is discussed more fully in the following, which describes example instructions that can be used to implement the present invention.

Resources are firstly reserved for a thread using a GETR instruction specifying the type of resource required, and can be freed again using a FREER instruction.

Ports can be used in input or output mode. In input mode a condition can be used to filter the data passed to the thread. A port can be used to generate events or interrupts when data becomes available as described below. This allows a thread to monitor several ports, only servicing those that are ready. Input and output instructions, IN and OUT, can then be used to transfer of data to and from ports once ready. In this case, the IN instruction inputs and zero-extends the n least significant bits from an n-bit port and the OUT instructions outputs the n least significant bits.

Two further instructions, INSHR and OUTSHR, optimise the transfer of data. The INSHR instruction shifts the contents of a register right by n bits, filling the left-most n bits with the data input from the n-bit port. The OUTSHR instruction outputs the n least significant bits of data to the n-bit port and shifts the contents of a register right by n bits.

OUTSHR	port < s[bits 0 for width(port)];	output from
port, s	s ← s >> width(port)	port and shift
INSHR port, s	s ← s >> width(d);	shift and input
	port > s[bits (bitsperword – width(d))	from port
	for width(d)]	

where the ">" represents an input and the "<" represents an output.

A port must be configured before it can be used. It is configured using the SETC instruction which is used to define

## 13

several independent settings of the port. Each of these has a default mode and need only be configured if a different mode is needed.

SETC port, mode port[ctrl] ← mode set port control

The effect of the SETC mode settings is described below. The first entry in each setting is the default mode.

Mode	Effect
OFF	port not active; pin(s) high impedance
ON	active
IN	port is an input
OUT	port is an output (but inputs return the current pin value)
EVENT	port will cause events
INTERRUPT	port will raise interrupts
DRIVE	pins are driven both high and low
PULLDOWN	pins pull down for 0 bits, are high impedance otherwise
PULLUP	pins pull up for 1 bits, but are high impedance otherwise
UNCOND	port always ready; inputs complete immediately
EQUAL	port ready when its value is equal to its DATA value
NE	port ready when its value is different from its DATA value
TRANSITION	port ready when its value changes towards its DATA value
GR	port ready when its value is greater than its DATA value
LS	port ready when its value is less than its DATA value

The DRIVE, PULLDOWN and PULLUP modes are only relevant when the port direction is OUT. The TRANSITION condition is only relevant for 1-bit ports and the GR and LS conditions are only relevant for ports with more than one bit.

Each port has a ready bit **37** which is used to control the flow of data through the port, and defines whether the port is able to complete input or output instructions. The ready bit is set in different ways depending on the port configuration. The ready bit is cleared when any of the SETC, SETD or SETV instructions are executed.

A port in input mode can be configured to perform conditional input. The condition filters the input data so that only data which meets the condition is returned to the program. When a condition is set, the IN and INSHR instructions will only complete when the port is ready. As described above, executing an input instruction on a port which is not ready will pause the thread. When ready, the port sets its ready bit which is signalled to the thread scheduler. The thread resumes and re-executes the input instruction. This time the port is ready, the data is returned and the ready bit **37** is cleared.

Once a port ready bit is set, the data value which satisfied the condition is captured so that the software gets the value which met the condition even if the value on the port has subsequently changed. When an IN or INSHR instruction is executed and the ready bit is set then the data is returned and the ready bit cleared. If the ready bit is not set then the thread is paused until the ready bit is set. If a condition is set then the data is compared against the condition and the ready bit is only set when the condition is met.

When the OUT or OUTSHR instruction is executed if the ready bit is clear then the data is taken by the port and the ready bit is set. If the ready bit is set then the thread is paused until it is cleared by the port.

In order to communicate between two threads, two channel ends need to be allocated, one for each thread. This is done using a GETR CHAN instruction. The two threads can then use the resource identifiers to transfer a data word using output and input instructions:

## 14

OUT	$d \leq s$
IN	$d \geq s$

If an output instruction is executed when the channel is too full to take the data then the thread which executed the instruction is paused. It is restarted when there is enough room in the channel for the instruction to successfully complete. Likewise, when an input instruction is executed and there is enough data available then the thread is paused and will be restarted when enough data becomes available. When it is no longer required, the channel can be freed using a FREER CHAN instruction. Otherwise it can be used for another message.

Events and interrupts allow resources (ports and channels) to automatically transfer control to a predefined event handler. The ability of a thread to accept events or interrupts is controlled by information held in the thread status register SR (see FIG. 4), and may be explicitly controlled using TSE and TSD instructions. This information comprises an event enable flag (EE) and an interrupt enable flag (IE).

TSE s	$SR \leftarrow SR \vee s$	thread state enable
TSD s	$SR \leftarrow SR \wedge \neg s$	thread state disable

The operand of these instructions should be one of:

EE	to enable or disable events
IE	to enable or disable interrupts

Events are handled in the same scope in which they were set up. Hence, on an event all the thread's state is valid, allowing the thread to respond rapidly to the event. The thread can perform input and output operations using the port which gave rise to an event whilst leaving some or all of the event information unchanged. This allows the thread to complete handling an event and immediately wait for another similar event.

The program location of the event handler must be set prior to enabling the event using the SETV instruction. Ports have conditions which determine when they will generate an event; these are set using the SETC and SETD instructions. Channels are considered ready as soon as they contain enough data or have room to accept data for output.

Event generation by a specific port or channel can be enabled using an event enable unconditional (EEU) instruction and disabled using an event disable unconditional (EDU) instruction. The event enable true (EET) instruction enables the event if its condition operand is true and disables it otherwise; conversely the event enable false (EEF) instruction enables the event if its condition operand is false, and disables it otherwise. These instructions are used to optimise the implementation of guarded inputs. Below are some example instruction formats for configuring events on ports, but it will be understood that the same instructions can apply in relation to channels.

SETV port, v	port[vector] ← v	set event vector
SETD port, d	port[data] ← d	set event data
SETC port, c	port[ctrl] ← c	set event control
EET port, b	port[enable] ← b; port[tid] ← thread	event enable true
EEF port, b	port[enable] ← $\neg b$ ; port[tid] ← thread	event enable false

## 15

-continued

EDU port	port[enable] ← false; port[tid] ← thread	event disable
EEU port	port[enable] ← true; port[tid] ← thread	event enable

Having enabled events on one or more resources, a thread can use a WAITEU instruction to wait for at least one event. This may result in an event taking place immediately with control being transferred to the event handler specified by the corresponding event vector with events disabled by clearing the EE (event enable) flag. Alternatively the thread may be suspended until an event takes place—in this case the EE flag will be cleared when the event takes place, and the thread resumes execution.

WAITET b	if b then SR[EE] ← true	event wait if true
WAITEF b	if ¬b then SR[EE] ← true	event wait if false
WAITEU	SR[EE] ← true	event wait
CLRE	SR[EE] ← false;	disable all events
	forall port	for thread
	if port[tid] = thread then	
	port[enable] ← false	

To optimise the common case of repeatedly waiting for one or more events until a condition occurs, conditional forms of the event wait instruction are provided. The WAITET instruction waits only if its condition operand is true, and the WAITEF waits only if its condition operand is false.

All of the events which have been enabled by a thread can be disabled using a single CLRE instruction. This disables event generation in all of the ports which have had events enabled by the thread. The CLRE instruction also clears the event-enabled status in the thread's status register.

In order to optimise the responsiveness of a thread to high priority resources, the TSE EE instruction can be used to enable events on a thread first before subsequently starting to enable the ports and/or channels and using one of the event wait instructions. This way, the processor can scan through the resources in priority order. This may cause an event to be handled immediately as soon as it is enabled.

In contrast to events, interrupts are not handled within the current scope and so the current PC and SR (and potentially also some or all of the other registers) must be saved prior to execution of the interrupt handler. On an interrupt generated by resource r the following occurs automatically:

```
SAVEPC ← PC;
SAVESR ← SR;
SR[EE] ← false;
SR[IE] ← false;
PC ← r[vector]
```

When the handler has completed, execution of the interrupted thread can be performed by an RFINT instruction.

RFINT	PC ← SAVEPC; return from interrupt
	SR ← SAVESR

An interrupt could interrupt a thread whilst suspended awaiting an event.

The following examples show how the instructions are used by threads to perform input, output and logical operations. In the examples, the following instructions are used:

LDFI: loads an instruction address into a register

LDI: loads a constant value into a register

EQI: produces a Boolean (truth) value if a register value equals a constant

## 16

OR: produces the logical OR of two register values

ADD: adds two register values

ADDI: adds a constant to a register value

SHL: shifts the contents of a register left

BBF: branches to another point in the program if a Boolean value is false

OUT: outputs data

The following shows example code for inputting an 8-bit byte serially from a pin. Each bit of data is input from a first port when a signal received at a second port from an external clock changes from 0 to 1 to indicate that the data should be taken. In a high level language, the operation looks like this:

```
15  PROC inbyte(clock, data)
    VAR byte;
    { FOR n = 0 FOR 8
      WHEN clock ? TO 1 DO { data ? bit ; byte = (byte << 1) + b ;
    }
    RETURN byte
20 }
```

The instruction level program for this is shown below.

```
25  SETC      clock, TRANSITION
    SETD      clock, 1
    LDI       byte, 0
    LDI       count, 0
    loop:
    IN         clock, c
    IN         data, bit
    SHL        byte, byte, data
    ADD        byte, byte, bit
    ADDI       count, count, 1
    EQI        go, count, 8
    BBF        go, Loop
```

It would be possible to execute two or more such code sequences at the same time by allocating each one of them to its own thread.

The following shows example code, using some of the above instructions, for implementing a NAND type process which wakes up whenever one of two inputs x and y changes state. The high level code is:

```
45  PROC nand(x, y, z) IS
    WHILE TRUE
    { WHEN x ? TO notx DO { notx := NOT notx; z ! (notx OR noty) }
    | WHEN y ? TO noty DO { noty := NOT noty; z ! (notx OR noty) }
    }
```

In low level code, the process comprises a single thread which initialises two ports x and y with vectors "xv" and "yv" respectively, and enables these ports to generate events. The corresponding instruction level program is as follows:

```
55  nand:
    LDI       notx, 1
    LDI       noty, 1
    SETC      x, TRANSITION
    SETD      x, notx
    SETC      y, TRANSITION
    SETD      y, noty
    LDFI      temp, xv
    SETV      temp, x
    LDFI      temp, yv
    SETV      temp, y
    EEU       x
60
65
```

17

-continued

---

```

EEU    y
WAIT
xv:
EQI     notx, notx, 0
SETD    x, notx
OR      temp, notx, noty
OUT     z, temp
WAIT
yv:
EQI     noty, noty, 0
SETD    y, noty
OR      temp, notx, noty
OUT     z, temp
WAIT

```

---

In operation, either the x-input changes or the y-input changes and control transfers either to xv or to yv. In either case, the response code executes five instructions, then waits for the next input state-change. Latency from input change to output change may be less than about 10 cycles. A 1 GHz processor can emulate 100 MHz logic.

As another example, the following shows a process for implementing D-type flip-flop logic which wakes up whenever an input changes state but only changes output when clocked by an external clock. The high-level program is:

---

```

PROC dtype(d, ck, q) IS
WHILE TRUE
{ WHEN d ? TO notd DO notd := NOT notd
| WHEN ck ? TO TRUE DO q ! NOT notd
}

```

---

The corresponding instruction level program is:

---

```

dtype:
IN      d, temp
EQI     temp, temp, 0
SETC    d, TRANSITION
SETD    d, temp
SETC    ck TRANSITION
LDI     temp, 1
SETD    ck, temp
LDFI    temp, dv
SETV    d, temp
LDFI    temp, ckv
SETV    ck, ckv
EEU     d
EEU     ck
WAIT
dv:
IN      d, temp
SETD    d, temp
WAIT
ckv:
EQI     temp, notd, 0
OUT     q, temp
WAIT

```

---

In operation, either the d-input changes or the ck-input changes. In either case, the response code executes three instructions, then waits for the next input state-change. Latency from input change to output change may be less than about 10 cycles. Again, a 1 GHz processor can emulate 100 MHz logic.

The following gives an example of some more complex

18

arrives (another way would be to only read the data on the clock, in which case there would be a non-zero hold time for the data). The output is calculated—in the example below by a lookup table—and output on the clock. A more complex function of the input could be calculated and this would potentially add more instructions at the point indicated below. However, notice that a processor can calculate some very complex functions (relative to a small LUT) in only a few instructions. The high-level code is:

---

```

PROC lookup(d, ck, q) IS
WHILE TRUE
{ WHEN d ? x : x != lastx DO lastx := x
| WHEN ck ? TO TRUE DO q ! lookup[lastx]
}

```

---

The corresponding instruction level program is:

---

```

logic:
IN      d, data
SETC    d, NE
SETD    d, data
SETC    ck, TRANSITION
LDI     temp, 1
SETD    ck, temp
LDFI    temp, dv
SETV    d, temp
LDFI    temp, ckv
SETV    ck, temp
EEU     d
EEU     ck
WAIT
dv:
IN      d, data
SETD    d, data
WAIT
ckv:
LDW     temp, lookup, data // replace this to change function
OUT     temp, q
WAIT

```

---

In operation, either the d-input changes or the ck-input changes. In either case, the response code executes three instructions, then waits for the next input state-change. Latency from input change to output change may be less than about 10 cycles. Again, a 1 GHz processor can emulate 100 MHz logic.

Note also that the above examples demonstrate how a given thread can handle multiple activities, such as multiple events.

In contrast to events, interrupts require state to be saved on entry to an interrupt handler and restored on exit in order to make registers available for use within the handler. In addition, the handler will normally need to retrieve state from when it was last entered and save it ready for when it is next entered. A simple example of an interrupt handler is shown below. This uses some additional instructions:

LDWSP loads a value from memory using the stack pointer  
 STWSP stores a value to memory using the stack pointer  
 LDWDP loads a value from memory using the data pointer  
 STWDP stores a value to memory using the data pointer  
 EXTSP used to extend the stack to make space for new values

LDAWSP used to discard values from the stack

This example inputs a byte of data one bit at a time; in contrast to the above example using events it uses an interrupt handler. The high level program is:

19

---

```

PORT clock: INT, TRANSITION, 1;
VAR byte;
VAR count;
byte := 0;
count := 0;
ENABLE clock;
// program to be interrupted follows here

```

---



---

```

HANDLER inbyte( )
{ data ? bit ; byte := (byte << 1) + bit;
  count := count + 1;
  IF count = 8 THEN DISABLE clock
}

```

---

When the port is enabled to generate interrupts, the interrupt handler is entered every time an external clock makes a transition to logic 1. The handler takes data bits and forms a byte. The byte, together with a count of bits input are stored in locations in memory and accessed via the data pointer. When 8 bits have been input, the handler disables further interrupts leaving the byte ready for a program to use. The corresponding instruction level program is:

---

```

SETD    clock, 1
SETC    clock, TRANSITION
SETC    clock, INT
LDI     r0, 0
STWDP   r0, byte
STWDP   r0, n
EEU     clock
... // program which will be interrupted
int: // interrupt handler

```

---



---

```

EXTSP    2
STWSP    R0, 0
STWSP    R1, 1
IN       clock, r0
IN       data, r0
LDWDP    r1, byte
SHL      byte, byte, 1
ADD      byte, byte, r0
STWDP    r1, byte
LDWDP    r0, n
ADDI     r0, r0, 1
STWDP    r0, n
EQI      r0, r0, 8
EEF      clock, r0
LDWSP    R1, 1
LDWSP    R0, 0
LDAWSP   2
RFINT

```

---

From the above description and examples, it can be seen how associating activity at respective ports with respective threads, and scheduling those threads based on events arising from that activity, advantageously provides a processor which can respond quickly to external stimuli.

It will be appreciated that the above embodiments are described only by way of example. In other embodiments, different sets of registers and instructions may be provided depending on the desired specifications of the chip. In some embodiments, thread identifiers need not be transmitted to ports but could remain the responsibility of the thread scheduler, or be stored elsewhere. Alternatively, each thread could be given an individual ready flag at the port, such that the thread identifier is passed to the port to select the correct ready signal but the thread identifier need not be returned to the

20

thread scheduler upon detection of the activity. Further, conditions and/or condition data need not be transmitted to ports. Instead conditions could be preconfigured at ports and/or conditions could be evaluated at the thread scheduler or elsewhere. Threads may be scheduled based on activity from other sources other than ports and channels. Different interconnects may be provided between the various components of the processor. Also, the invention is not specific to use in a mobile terminal with a mobile applications processor. Other applications and configurations will be apparent to the person skilled in the art. The scope of the invention is not limited by the described embodiments, but only by the following claims.

The invention claimed is:

1. A processor comprising:  
an execution unit;

a thread scheduler configured to schedule a plurality of threads for execution by the execution unit in dependence on a respective runnable status for each respective thread, each thread represented in a respective one of a plurality of sets of thread registers; and

a plurality of event sources being hardware resources comprising one or more ports for transferring data between the processor and a device external to the processor, and/or one or more channels of an interconnect system for communicating between one of said plurality of sets of thread registers representing one of said plurality of threads and another set of thread registers representing another thread;

wherein the execution unit is configured to execute thread scheduling instructions being instructions of an instruction set of the execution unit which manage said runnable statuses;

wherein the thread scheduling instructions of the instruction set include a source event enable instruction, the execution unit being configured to be operated by an opcode of the event enable instruction to send a signal to a specified one of said ports or channels via a direct hardwired connection between the execution unit and the specified port or channel in order to set the specified port or channel to a mode in which it generates a respective event dependent on activity occurring at the specified port or channel, the direct hardwired connection being a connection not via a bus;

wherein each of said ports or channels is operable to be associated with one of the plurality of threads and with a specified one of a plurality of different potential code locations; and

wherein the thread scheduling instructions also include a wait instruction which sets one of said runnable statuses to suspend the respective thread until an associated one of said ports or channels generates its respective event, such that execution of the respective thread continues from the specified one of the plurality of different potential code locations.

2. The processor according to claim 1, wherein the thread scheduling instructions include a source event disable instruction, the execution unit being configured to be operated by an opcode of the event enable instruction to send a signal from the execution unit to the specified port or channel via said direct hardwired connection in order to set the specified port or channel to a mode in which it does not generate the respective event.

3. The processor according to claim 1, wherein the thread scheduling instructions include: a thread event enable instruction which sets a control status of a thread to event-enabled to allow the event-enabled thread to accept events, and a thread

21

event disable instruction which sets a control status of a thread to event-disabled to stop the event-disabled thread from accepting events.

4. The processor according to claim 3, wherein each of said sets of thread registers comprises a register for storing said control status.

5. The processor according to claim 3, wherein the execution unit is arranged to execute said thread event enable instruction and subsequently execute a plurality of said source event enable instructions each for a respective port or channel.

6. The processor according to claim 1, wherein the execution unit is configured to execute one or more additional instructions each of which configures a parameter associated with a specified one of said ports or channels wherein each of said ports or channels comprises one or more registers for storing the parameter being configured by each of said one or more additional instructions.

7. The processor according to claim 6, wherein the one or more additional instructions include at least a set condition instruction which sets a condition upon which a specified one of said ports or channels generates the respective event.

8. The processor according to claim 7, wherein the one or more additional instructions include a set data instruction which provides data associated with the condition.

9. The processor according to claim 6, wherein the one or more additional instructions include at least a set vector instruction which sets a continuation point vector specifying said one of the plurality of potential code locations associated with the specified event source.

10. The processor according to claim 1, wherein the thread scheduling instructions further include an interrupt enable instruction which sets a control status of a thread to interrupt-enabled to allow the interrupt-enabled thread to accept interrupts, and an interrupt disable instruction which sets a control status of a thread to interrupt-disabled to stop the interrupt-disabled thread from accepting interrupts.

11. The processor according to claim 1, wherein the thread scheduling instructions further include: an input instruction which pauses a thread pending input of data from one of said ports or channels, and an output instruction which pauses a thread pending the availability of one of said ports or channels for outputting data; wherein continued execution of a thread paused by said input and output instructions resumes where execution was paused.

12. The processor according to claim 11, wherein the processor is arranged to execute a first one of said source event enable instructions to enable a first port or channel to generate a first event dependent on a first activity occurring at the first port or channel, to execute at least one of said additional instructions to configure one or more parameters for the first port or channel, to execute said wait instruction to suspend the respective thread pending the first activity, to disable the first port or channel from generating the first event after occurrence of the first activity, to leave the one or more parameters configured for the first port or channel after said disabling of the first port or channel, and to re-use said one or more parameters by executing a second of said source event enable instructions to enable the first port or channel to generate a further instance of the first event dependent on a further occurrence of the first activity at the first port or channel.

13. The processor according to claim 12, wherein the processor is arranged to execute at least one of said input and output instructions between said execution of said wait instruction and said execution of said second event source enable instruction.

22

14. The processor according to claim 1, wherein the thread scheduling instructions include a clear instruction to disable all events for a thread.

15. The processor according to claim 1, wherein the port or channel specified by at least one of the one or more source event enable instructions comprises the channel, such that said activity comprises activity occurring over the channel.

16. The processor according to claim 1, wherein the port or channel specified by at least one of the one or more source event enable instructions comprises the port, such that said activity comprises activity occurring at the port.

17. The processor according to claim 1, wherein the execution unit is adapted to execute a plurality of said source event enable instructions each of which sets a respective port or channel to a mode in which it generates an event dependent on activity occurring at that port or channel, and said wait instruction sets said one of the runnable statuses to suspend the respective thread until generation of an event from any of the respective ports or channels.

18. The processor according to claim 1, wherein said wait instruction can be either a wait enable true instruction which waits only if its condition operand is true, or a wait enable false instruction which waits only if its condition operand is false.

19. The processor according to claim 1, wherein said source event enable instruction can be either a source event enable true instruction which enables the specified port or channel to generate events if an operand is true and disables it otherwise, or a source event enable false instruction which enables the specified port or channel to generate events if an operand is false and disables it otherwise.

20. The processor according to claim 1, wherein the processor is adapted to automatically disable a thread from accepting events upon occurrence of the respective event without executing a thread event disable instruction.

21. The processor according to claim 1, wherein the processor is adapted to automatically disable the port or channel from generating events upon occurrence of a respective event without executing a source event disable instruction.

22. The processor according to claim 1, wherein the processor is adapted to complete the wait instruction and continue the respective thread immediately if the respective event is generated on or before execution of the wait instruction.

23. The processor of claim 1, wherein said device external to the processor is an off-chip device on another chip than said processor.

24. The processor of claim 1, wherein said device external to the processor comprises another processor.

25. The processor of claim 1, wherein the association of ports or channels with one of the plurality of threads comprises storing a thread identifier in a thread identifier register of the port or channel.

26. A method of controlling a thread scheduler to schedule a plurality of threads for execution by an execution unit within a processor, each thread represented in a respective one of a plurality of sets of thread registers; and the processor comprising a plurality of event sources being hardware resources comprising one or more ports for transferring data between the processor and a device external to the processor, and/or one or more channels of an interconnect system for communicating between one of said plurality of sets of thread registers representing one of said plurality of threads and another set of thread registers representing another thread; and the method comprising:

scheduling the plurality of threads in dependence on a respective status for each respective thread; and

23

operating the execution unit to execute thread scheduling instructions being instructions of an instruction set of the execution unit for managing statuses of threads;

wherein said thread scheduling instructions include a source event enable instruction having an opcode which operates the execution unit to send a signal to a specified one of said ports or channels via a direct hardwired connection between the execution unit and the specified port or channel in order to set the specified port or channel to a mode in which it generates a respective event dependent on activity occurring at the specified port or channel, the direct hardwired connection being a connection not via a bus;

wherein each of said ports or channels is operable to be associated with one of said threads and with a specified one of a plurality of different potential code locations; and

wherein the thread scheduling instructions also include a wait instruction which sets one of said runnable statuses to suspend the respective thread until an associated one of said ports or channels generates its respective event, such that execution of the respective thread continues from the specified one of the plurality of different potential code locations.

27. The method according to claim 26, wherein the thread scheduling instructions include a source event disable instruction having an opcode which operates the execution unit to send a signal to the specified port or channel via said direct hardwired connection in order to set the specified port or channel to a mode in which it does not generate the respective event.

28. The method according to claim 26, wherein the thread scheduling instructions include: a thread event enable instruction which sets a control status for a thread to event-enabled to allow the event-enabled thread to accept events, and a thread event disable instruction which sets a control status of a thread to event-disabled to stop the event-disabled thread from accepting events.

29. The method according to claim 28, comprising storing said control status in at least one register of each of said sets of thread registers.

30. The method according to claim 28, comprising executing said thread event enable instruction and subsequently executing a plurality of said source event enable instructions each for a respective source.

31. The method according to claim 26, further comprising executing one or more additional instructions each of which configures a parameter associated with a specified one of said ports or channels, and storing, in one or more registers of the specified port or channel, the parameter being configured by each of the one or more additional instructions.

32. The method according to claim 31, wherein said one or more additional instructions include at least a set condition instruction which sets a condition upon which a specified one of said ports or channels generates the respective event.

33. The method according to claim 32, wherein said one or more additional instructions include a set data instruction which provides data associated with the condition.

34. The method according to claim 31, wherein said one or more additional instructions include at least a set vector instruction which sets a continuation point vector specifying said one of the plurality of different potential code locations associated with the specified port or channel.

35. The method according to claim 26, wherein the thread scheduling instructions further include: an interrupt enable instruction which sets a control status for a thread to interrupt-enabled to allow the interrupt-enabled thread to accept inter-

24

rupts, and an interrupt disable instruction which sets a control status for a thread to interrupt-disabled to stop the interrupt-disabled thread from accepting interrupts.

36. The method according to claim 26, wherein the thread scheduling instructions further include: an input instruction which pauses a thread pending input of data from one of said ports or channels, and an output instruction which pauses a thread pending the availability of one of said ports or channels for outputting data; wherein continued execution of a thread paused by said input and output instructions resumes where execution was paused.

37. The method according to claim 36, comprising:

executing a first one of said source event enable instructions to enable a first port or channel to generate a first event dependent on a first activity occurring at the first port or channel,

executing at least one of said additional instructions to configure one or more parameters for the first event source,

executing said wait instruction to suspend the respective thread pending the first activity,

disabling the first port or channel from generating the first event source after occurrence of the first activity,

leaving said one or more parameters configured for the first port or channel after said disabling of the first port or channel, and

re-using said parameters by executing a second of said source event enable instructions to enable the first port or channel to generate a further instance of the first event dependent on a further occurrence of the first activity at the first port or channel.

38. The method according to claim 37, comprising executing at least one of said input and output instructions between said execution of said wait instruction and said execution of said second event source enable instruction.

39. The method according to claim 26, wherein said thread scheduling instructions include a clear instruction to disable all events for a thread.

40. The method according to claim 26, wherein the port or channel specified by at least one of the one or more source event enable instructions comprises the channel, such that said activity comprises activity occurring over the channel.

41. The method according to claim 26, wherein the port or channel specified by at least one of the one or more source event enable instructions comprises the port, such that said activity comprises activity occurring at the port.

42. The method according to claim 26, comprising:

executing a plurality of said source event enable instructions to set each of a plurality of respective ports or channels to a mode in which it generates an event dependent on activity occurring at that port or channel, and

executing said wait instruction to set said one of the runnable statuses to suspend the respective thread until generation of an event from any of the respective ports or channels.

43. The method according to claim 26, wherein said wait instruction can be either a wait enable true instruction which waits only if its condition operand is true, or a wait enable false instruction which waits only if its condition operand is false.

44. The method according to claim 26, wherein said source event enable instruction can be either a source event enable true instruction which enables the specified port or channel to generate events if an operand is true and disables it otherwise, or a source event enable false instruction which enables the specified port or channel source to generate events if an operand is false and disables it otherwise.

## 25

45. The method according to claim 26, comprising automatically disabling a thread from accepting events upon occurrence of the respective event without executing a thread event disable instruction.

46. The method according to claim 26, comprising automatically disabling the event source from generating events upon occurrence of a respective event without executing a source event disable instruction.

47. The method according to claim 26, comprising completing the wait instruction and continuing the respective thread immediately if the respective event is generated on or before execution of the wait instruction.

48. A processor comprising:

an execution unit;

a thread scheduler configured to schedule a plurality of threads for execution by the execution unit in dependence on a respective runnable status for each thread; and

a plurality of event sources for generating events;

wherein the execution unit is configured to execute thread scheduling instructions which manage said runnable statuses, the thread scheduling instructions including one or more source event enable instructions each of

## 26

which sets a specified one of said event sources to a mode in which it generates one of said events dependent on activity occurring at the specified source, and a wait instruction which sets one of said runnable statuses to suspend the respective thread until generation of one of said events upon which continued execution of that thread depends;

wherein a continuation point vector for the respective thread is associated with the specified event source, and, if suspended by said wait instruction, said continued execution comprises retrieving the continuation point vector for the respective thread and continuing execution from the continuation point vector; and

wherein the thread scheduling instructions further include an input instruction which pauses a thread pending input of data from an event source, and an output instruction which pauses a thread pending the availability of an event source for outputting data, wherein continued execution of a thread does not involve retrieval of a continuation point vector for that thread if paused by said input and output instructions instead of the wait instruction.

\* \* \* \* \*